NATIONAL AND REGIONAL IMPLICATIONS OF INTERNET DATA CENTER GROWTH

Jennifer Mitchell-Jackson[a], Jonathan G. Koomey[b*], Michele Blazek[c], Bruce Nordman[b]

[a]*Energy and Resources Group, 310 Barrows Hall, University of California, Berkeley, CA 94720*

[b]*Lawrence Berkeley National Laboratory 1 Cyclotron Road, Mail stop 90-4000, Berkeley, CA 94720*

[c]*AT&T, Suite 3188, 4430 Rosewood Drive, Pleasanton, CA 94588*

*Corresponding author. Tel.: +1-510-486-5974; fax: +1-510-486-4247.

Email address: JGKoomey@lbl.gov (J. Koomey)

ABSTRACT

The electricity consumption of Internet data centers is a growing concern to utility demand forecasters, data center facility managers, energy analysts and policy makers. Using estimates of U.S. computer room floor space and total computer room power density, we found that U.S. data centers in the aggregate require less than 500 MW of power, and use only a fraction of a percent (less than 0.12% in 2000) of the electricity consumed nationwide. In this paper, our order-of-magnitude estimate suggests that energy demands of Internet data centers do not represent an enormous new burden on the electricity industry as a whole. The fact that these facilities tend to be concentrated in certain areas, however, may mean that there will be significant regional electricity demands in some parts of the country.

Keywords: *data centers, power, electricity consumption, power densities*

INTRODUCTION

Energy planning in the United States is highly dependent on forecasted demand, yet for a growing number of Internet data centers, power requirements are greatly exaggerated. These exaggerated estimates can lead to the building of excess generation, transmission, and distribution capacity. The goal of this paper is to provide a better understanding of the electricity demands of this industry in order to establish realistic growth projections for the future. Below, we estimate the aggregate electricity demand from U.S. data center hosting facilities based on the total computer room floor space and the average power density used to support this space.

ESTIMATING COMPUTER ROOM FLOOR SPACE IN U.S. DATA CENTERS

Only a few rough estimates of total U.S. data center floor space exist; and even fewer are truly indicative of Internet growth. Estimates of total building floor area for data center facilities include significant amounts of space that is not critical to a data center's main function. In a data center facility, 15% to 75% of the building's area is usually designated for meeting rooms, offices, restrooms and hallways. None of this space accurately reflects the growth of this industry. Estimates of computer room floor area, or the space devoted to the computer equipment, are most representative of this growth. Computer room floor area includes the area under the racks of computer equipment as well as the electrically inactive areas, such as the aisles between racks of computer equipment. It does not include office space, lobbies, bathrooms, space set aside for potential future expansion, or mechanical equipment rooms outside of the computer rooms.

There are at least two estimates of computer room floor area. Both of these estimates are based on surveys of large data center hosting companies. One estimate by Salomon Smith Barney anticipated that the there would be nearly 900,000 square meters of data center computer rooms built or under construction in the United States at the end of 2000, and that this number would rise to approximately 1.7 million square meters by the end of 2001 (Mahedy et al. 2000). A second estimate by the Yankee Group predicted that there would be 860,000 square meters of computer rooms at the end of 2000—approximately 6% less than in the Salomon Smith Barney report (Yankee Group 2000).

Based on the Yankee Groups' numbers for 1998 through 2000, and using the growth rate from the Salomon Smith Barney report to make an estimate of the computer room floor area in 2001, total computer room floor area in U.S. data centers is estimated to have jumped from less than 100,000 square meters in 1998 to roughly 900,000 square meters at the end of 2000 (Yankee Group 2000).[1] (See Figure 1.) Current projections indicate that there will be approximately 1.6 million square meters by the end of 2001 and that the amount of floor area will continue to grow, although the recent economic events in this industry have led to a significant slowdown of data center construction.[2]

Although these are currently the best available estimates, the numbers shown in Figure 1 are still rough. Neither of the estimates mentioned above included all of the companies that own data centers, thus the aggregate area may be underestimated due to the exclusion of some facilities. It is also possible that these estimates are overstated. The Salomon Smith Barney report, for example, indicates that at the end of 2000, HostPro, Inc. had five data centers with approximately 5,950 square meters of computer room space, as well as 8,730 square meters of data center computer room space in "unspecified" locations (Mahedy et al. 2000). According to the HostPro website, however, the company currently has only five data centers.[3] It is possible

[1] Note that this number includes only the computer room floor area and not the floor area of the entire building.

[2] 2001 value estimated from the growth rate in Mahedy et al., 3 August 2000.

[3] Information from the HostPro website, www.hostpro.com, viewed April 12, 2001.

that HostPro intended to build additional facilities, but has not yet done so. It should be noted, however, that the HostPro estimate included in the Salomon Smith Barney report might be more than twice as large as HostPro's actual computer room floor area.

ESTIMATING AVERAGE POWER DENSITY IN DATA CENTER COMPUTER ROOMS

In detailed work on the energy consumption of data centers, we found that there is much confusion over the power requirements at data center facilities (Mitchell-Jackson 2001, Mitchell-Jackson et al. 2001). We identified many reasons for the overstated energy requests such as erroneously extrapolating total building power requirements from the power density of an isolated area or using design rather than actual power requirements.[4] In addition, many estimates neglect to include the power necessary for end uses such as heating, ventilation, and cooling (HVAC) or other equipment that is necessary to the functioning of the data center.

Because of the confusion over this issue, in our earlier work we developed a common metric, *total computer room power density*, which can be used to estimate the power requirements of this industry (Mitchell-Jackson et al. 2001). For the purposes of this article, we define total computer room power density as the power drawn by the computer equipment and all of the supporting equipment such as power distribution units, uninterruptible power supplies, HVAC and lights (in watts) divided by the computer room floor area (in square meters). While the total computer room power density is a difficult number to grasp, it is a meaningful indicator of power needs and can be compared between buildings of different sizes as well as between data centers at different stages of development.

To determine an average value for total computer room power density, we reviewed billing data for five data center facilities throughout the country. (See Table 1.) Electricity billing data were used to find average demand in the month with the highest consumption. (This was usually the most recent month.) The highest average power demand for the facility was then divided by the computer room floor area. This estimate of total computer room power density is an overestimate because it assumes that all of the power for the entire facility is used for the computer room; thus our estimate is an upper limit for the total computer room power density. Even these overestimates, however, indicate that the total computer room power density was always less than 430 W/m^2 (40 W/ft^2).

NATIONAL IMPACTS

Assuming that nationwide, approximately 900,000 square meters of data center space was devoted to computer equipment in 2000 and using a high-end average estimate of total computer room power density, across the country, these facilities would require less than 500 MW of power and would use a fraction of a percent of all electricity used nationwide. See Table 2.

Table 2 also shows that even with high estimates, by 2003, data center power requests will still equal less than one percent of all electricity consumption nationwide. Moreover, it is important to note that some unknown portion of this demand is not actually *new* electricity demand. Some of the computers in these data centers are just relocated from corporate office buildings to data center hosting facilities. In fact, a major market for some smaller data centers is convincing small businesses to move their Internet equipment out of a closet and into the safe and secure

[4] See Mitchell-Jackson et al. (Mitchell-Jackson 2001) for details on why data center power is often exaggerated.

environment of a hosting facility. In these situations, the previously dispersed computer equipment is often moved less than 15 kilometers from the office where it was originally located. Thus, this portion of data center electricity use would not represent new electricity demand.

Note that our numbers do not include corporate data centers, but “data center power use” within the commercial sector is difficult to capture because within corporations, there is a continuum from single isolated servers in an office, to groups of servers, to large rooms constructed specifically for computers. A “total corporate data center power use” value would be both difficult to capture, and not as meaningful as total commercial energy use. Kawamoto et al. (Kawamoto et al. 2001), however, estimate that all office, telecommunications, and network equipment in the United States use only 3% of the country’s total electricity consumption. Since corporate data center electricity use is a subset of this, the Kawamoto et al. estimate provides an upper bound for this subset.

REGIONAL IMPACTS

While data centers represent only a fraction of the total electricity consumed in the United States, regional electricity demands may be more significant. Despite the decentralized nature of the Internet, the communication system involved with Internet transactions has led to a concentration of electricity demands at locations along the nation’s fiber optic backbones. According to a Robertson Stephens report, there were approximately 320 data center hosting facilities in the United States at the end of 2000 (Juarez et al. 2001).[5] The majority of these hosting facilities are clustered in the major metropolitan areas shown in Figure 2.

Silicon Valley, California (consisting of four Bay Area counties: Alameda, San Francisco, San Mateo, and Santa Clara) is one of the largest data center hubs in the country. The report by Robertson Stephens estimates that there are 54 data center hosting facilities in this region (Juarez et al. 2001). According to this report, the hosting facilities in the Bay Area make up approximately 17% of all major hosting facilities in the United States. The Salomon Smith Barney report (which pre-dates the Robertson Stephens report and covers only 40 companies) agrees that approximately 15% of the data centers are being built in this region, but indicates that these data centers may represent less than 15% of the country’s data center computer room floor area, thus many Silicon Valley data centers may be smaller than average. This report also indicates that computer room floor space in the New York City area may rival the amount of floor space in the Bay Area (Mahedy et al. 2000).

In the Bay Area, estimates indicate that there were 93,000 square meters of data center space at the end of 2000 (Mahedy et al. 2000).[6] Assuming a power density of 538 W/m^2 (50 W/ft^2), data centers could require 50 MW of power in the Bay Area alone—approximately 10% of the total demanded by data centers nationwide. In the Bay Area, this could mean approximately 438 GWh of electricity a year, or roughly 1.2% of electricity consumption in this area (Mahedy et al. 2000).

---

[5] This was about 45 more than Salomon Smith Barney’s projections (Mahedy et al. 2000) but the Robertson Stephens report (Juarez et al. 2001) covered 60 companies while the Salomon Smith Barney report surveyed only 40.

[6] Note that this estimate does not include the 2.2 million (gross) square foot U.S. Dataport facility planned for the area.

Because of the clustering of these facilities, it is possible that local power problems will occur in regions where data centers are concentrated. California's energy crisis, however, is not the result of Internet data center growth. In a recent *Computer World* article that power demands "skyrocketed by 12%" in the heart of Silicon Valley compared to a statewide growth rate of 2% or 3% and the implication that this is the result of Internet growth, the California Energy Commission's webpage on "Silicon Valley Electricity Consumption" (seen in Table 3) shows that total electricity use in the Silicon Valley has not grown at a substantially higher rate than in the rest of California (CEC 2001, Hall 2001). Both Silicon Valley and statewide electricity consumption grew by roughly 16% over this 11 year period.

While not responsible for California's power crisis, utilities may, however, experience power problems in areas where data centers are concentrated. Data centers appear to be ideal electricity customers: they demand a relatively steady amount of power 24 hours a day. In reality, however, even after these facilities are built, utilities do not always have an accurate estimate of demand since this is a new industry and changes are occurring rapidly. For a utility, not having a good sense of the data center's electricity demand leads to difficulties in providing supply and hedging against risks. Utilities also face the challenge of meeting the customer's demand for infrastructure. This is especially difficult given the distinct differences in timing and planning cycles between utilities and the Internet industry. Utilities are accustomed to getting two or three years notice for new large office buildings and industrial centers. Now, they face the challenges of putting in power lines, transformers, and substations within a few months.

Faced with new requests for power, utilities must decide how to respond. Building new infrastructure and acquiring new power resources to meet demand will have serious costs. If priced correctly, these costs can be recouped through electricity charges, but if demand is significantly overstated, the utilities will spend more on infrastructure than they could ever hope to recover. Furthermore, utilities may build infrastructure to meet the power demands of an industry that could fade as rapidly as it appeared. In order to avoid excessive risk, several utilities have started to charge data centers for electricity infrastructure based on their initial requests for power. The utilities then pay the developers' back in portions over depending on how a site's electricity usage progresses towards the original load estimate (Ahlberg 2001, Cook 2000). To utilities, accurate estimates of power needs are extremely valuable.

THE POTENTIAL FOR COMBINED HEAT AND POWER IN DATA CENTERS

One policy question that has arisen recently is the potential for combined heat and power (CHP) technologies to supply cooling and electricity to data centers. In many ways, these facilities represent ideal loads for CHP, in that they run every hour of the year, with relatively small variations in electric loads and cooling needs. Table 4 shows a simple and illustrative calculation based on the data presented above. The calculation begins with total electricity used by data centers from Table 2, and estimates the contribution of different end-uses to that total, based on the characteristics of the data center floor area for the Bay Area data center explored in detail in (Mitchell-Jackson et al. 2001). We then calculate the cooling load as the amount of heat from computers, lights, and auxiliary equipment, and use a typical Coefficient of Performance (COP) for absorption cooling to determine the amount of heat needed from the CHP unit to remove the heat represented by the cooling load from the building. Finally, we multiply the heat needed to supply absorption cooling by the electricity production rate (EPR) for a typical gas-fired internal combustion (IC) engine cogeneration unit. The EPR characterizes the amount of electricity generated from a CHP plant divided by the amount of useful heat from that same plant (for details on definitions and performance characteristics of typical CHP systems, see (Krause et al. 1994)).

Table 4 shows that a significant potential for CHP exists for U.S. data centers. Total potential CHP generation in 2000 is about 5 TWh/year, which corresponds to about 0.6 GW of capacity. The range of generation for the projected floor area and power needs in 2003 is from 7 to 26 TWh, with the middle estimate totaling about 15 TWh/year. The electricity generated by the CHP is more than enough to supply the non-chiller loads in these data centers, so these facilities have the potential to become net producers of electricity. Of course, local air pollution concerns may influence whether such installations are feasible or desirable. Fuel cells may mitigate some concerns about local air pollution, but more research is needed to determine the cost and operating characteristics of this technology in this application.

CONCLUSIONS

Based on our findings for total computer room power density and the estimate of computer room floor area for 2000, data center facilities in the aggregate require less than 500 MW of power, and use a fraction of a percent (less than 0.12% in 2000) of all electricity used nationwide. This is less than one-tenth of the electricity consumed by the U.S. iron and steel industry (Margolis and Brindle 2000). Moreover, some unknown portion of this demand is not actually *new* electricity demand. Some of the computers in these data centers are just relocated from corporate office buildings to data center hosting facilities. The electricity use by U.S. data centers, therefore, does not represent an enormous new burden on the energy industry, although there may be regional power constraints in areas where data center facilities are concentrated.

Current estimates need to be revised to reflect real power demands. Additional detailed studies of current energy requirements, as well as better estimates of total U.S. computer room floor area, will help all parties to better understand the needs of this industry.

Acknowledgements

We would like to acknowledge the support of the Environmental Protection Agency, Office of Air and Radiation, and the comments of Ernst Worrell, Skip Laitner, and anonymous journal reviewers.

This work was supported by the Office of Air and Radiation of the U.S. Environmental Protection Agency. Prepared for the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

References

Ahlberg, E. 2001. "Electricity Utilities Fear Drought at Server Farms." *Dow Jones Newswires.* January 30. <http://www.djnewswires.com/>.

CEC. 2001. *Web site on Silicon Valley Electricity Consumption*. California Energy Commission. <http://www.energy.ca.gov/silicon_valley_consumption.html>.

Cook, J. 2000. "Internet data gain is a major power drain on local utilities." *Seattle Post Intelligencer*. September 5. <Read on the web at http://seattlep-i.nwsource.com/>.

Hall, Mark. 2001. "Net Blamed as Crisis Roils California." In *Computer World.* January 15. <Read on the web at http://www.computerworld.com/cwi/story/0,1199,NAV47_STO56341,00.html>.

Juarez, Richard A., Michael T. Alic, Karkhaniz Chetan S., and Brett D. Johnson. 2001. *Space Dex III—Hosting Space: Not All Space Is Created Equal—Smart, Complex Space Takes Center Stage*. Boston, MA: Robertson Stephens Inc. January 29. <http://www.rsco.com>.

Kawamoto, Kaoru, Jonathan Koomey, Bruce Nordman, Richard E. Brown, Maryann Piette, Michael Ting, and Alan Meier. 2001. "Electricity Used by Office Equipment and Network Equipment in the U.S." *Forthcoming in Energy--The International Journal (also LBNL-45917).*

Krause, Florentin, Jonathan Koomey, Hans Becht, David Olivier, Giuseppe Onufrio, and Pierre Radanne. 1994. *Fossil Generation: The Cost and Potential of Low-Carbon Resource Options in Western Europe*. **Energy Policy in the Greenhouse**. Volume II, Part 3C. El Cerrito, CA: International Project for Sustainable Energy Paths.

Mahedy, Stephen, Dan Cummins, and Danny Joe. 2000. *Internet Data Centers: If Built…Will They Come?* New York, NY: Salomon Smith Barney. August 3.

Margolis, Nancy, and Ross Brindle. 2000. *Energy and Environmental Profile of the U.S. Iron and Steel Industry*. Washington, DC: Prepared by Energetics Inc. for the U.S. Department of Energy Office of Industrial Technologies. DOE/EE-0229. August.

Mitchell-Jackson, Jennifer. 2001. *Energy Needs in an Internet Economy: A Closer Look at Data Centers*. M.S. Thesis, Energy and Resources Group, University of California, Berkeley.

Mitchell-Jackson, Jennifer, Jonathan Koomey, Bruce Nordman, and Michele Blazek. 2001. "Data Center Power Requirements: Measurements From Silicon Valley." *Submitted to Energy--the International Journal (also LBNL-48554)*. July.

US DOE. 2000. *Annual Energy Outlook 2001, with Projections to 2020*. Washington, DC: Energy Information Administration, U.S. Department of Energy. DOE/EIA-0383(2001). December.

Yankee Group. 2000. *Executive Summary of The U.S. Collocation Market: High-Tech Real Estate Heats Up*. Boston, MA: Yankee Group. Viewed on the web, March 28, 2001. <www.yankeegroup.com>.

Location	Units	Data Center A	Data Center B	Data Center C	Data Center D	Data Center E
Building Area	m^2	11,643	10,684	14,321	NA	33,292
Computer Room Area	m^2	2,555	3,716	4,181	4,476	3,577
Building Power Density	W/m^2	118	31	106	NA	40
Upper Limit for Total Computer Room Power Density	W/m^2 (W/ft^2)	355 (33)	88 (8)	363 (34)	412 (38)	376 (35)

Table 1. Comparison of power densities at five U.S. data centers. The Upper Limit for Total Computer Room Power Density for data centers B-E was calculated by dividing the average power demand for the entire facility (from billing data) by the computer room area. This number includes all of the power used by the entire building and is therefore an overestimate. The estimate for data center A is based on a more detailed study of power requirements (see reference 4).

	Units	2000	2003 low	2003 mid	2003 high
Computer Room Floor Area	Million m^2 (Million ft^2)	0.9 (9.5)	1.9 (20)	2.3 (25)	2.8 (30)
Total computer room power density	W/m^2 (W/ft^2)	538 (50)	377 (35)	646 (60)	915 (85)
Data center total power	MW	475	700	1500	2550
Data center electricity use	TWh	4	6	13	22
US electricity use	TWh	3364	3608	3608	3608
Data centers as % total use	***%***	***0.12%***	***0.17%***	***0.36%***	***0.62%***

[a] Total U.S. electricity use from Energy Information Energy Outlook Administration's Annual 2001 (US DOE 2000).

Table 2. Nationwide electricity demands for data center hosting facilities.

Year	**Silicon Valley Total Electricity Consumption**		**Statewide Total Electricity Consumption**	
	in million kWh	% growth	in million kWh	% growth
1990	31,436		228,038	
1991	31,140	-1%	222,260	-3%
1992	31,587	1%	226,988	2%
1993	31,679	0%	227,624	0%
1994	31,467	-1%	230,097	1%
1995	32,220	2%	230,990	0%
1996	32,911	2%	239,168	4%
1997	34,469	5%	247,437	3%
1998	34,289	-1%	244,510	-1%
1999	35,360	3%	252,800	3%
2000	36,616	4%	264,429	5%
	Overall Growth 1990-2000	16%	**Overall Growth 1990-2000**	16%
	Average Annual Growth	1.6%	**Average Annual Growth**	1.5%

Table 3. Electricity consumption in Silicon Valley versus the state of California. Source: California Energy Commission (CEC 2001).

		2000	2003 Low	2003 Mid	2003 High
Total electricity used by data centers (1)	TWh.e	4.2	6.1	13.1	22.3
Computer room, including lights and aux equipment	TWh.e	2.6	3.8	8.2	13.9
Central chiller plant	TWh.e	0.6	0.9	1.9	3.2
Fans, CRAC units, other ventilation	TWh.e	1.0	1.4	3.1	5.3
Total cooling load	TWh.th	2.6	3.8	8.2	13.9
Total heat needed to meet cooling load using absorption cooling (2, 5)	TWh.th	3.7	5.4	11.6	19.8
Total electricity generated from CHP (3, 4, 6)	TWh.e	4.9	7.2	15.4	26.2
% of non-chiller electricity generated by CHP		*137%*	*137%*	*137%*	*137%*
Total capacity (gas industrial IC CHP)	GW.e	0.6	0.9	1.9	3.3

Note: TWh.e = Terawatt-hours of electricity; TWh.e = Terawatt-hours of thermal energy (heat); GW.e = Gigawatts of electricity generation capacity.
(1) Total electricity used by data centers taken from Table 2. Breakdown by end use extrapolated from detailed work on electricity consumption in one Bay Area data center from Mitchell-Jackson et al. (2001).
(2) COP of absorption cooling =cooling out/heat in = 0.7. Assumes indirect-fired single-effect absorption cooling.
(3) Electricity production rate = electricity out/useful heat out, for a 25 MW industrial Internal Combustion (IC) engine = 1.32. From Krause et al. (1994), TLF, p. A.10.7.7.
(4) Capacity factor (industrial IC CHP) = Equivalent availability = 92%. From Krause et al. (1994), p. A.10.7.7.
(5) Total heat needed to meet cooling load = Cooling load/COP.
(6) Total electricity generated = total heat needed to meet cooling load x electricity production rate. Assumes a gas fired industrial internal combustion engine, 25 MW electrical capacity, operated in thermal load following mode, taken from Krause et al. (1994).

Table 4: Total Combined Heat and Power (CHP) potential to meet data center cooling loads and generate electricity

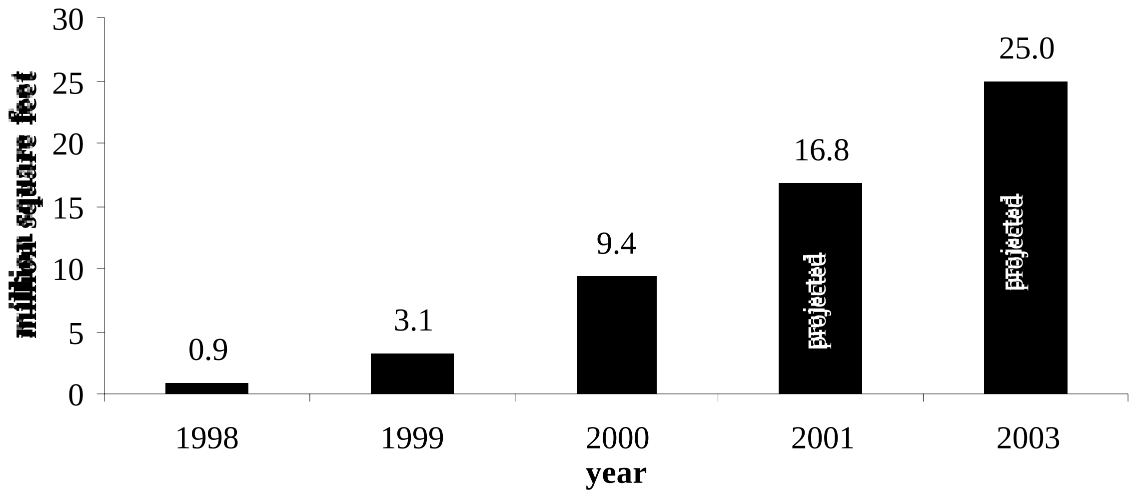

Figure 1. Computer room floor space in U.S. data centers. Source: Yankee Group, 2000 (Yankee Group 2000); 2001 value estimated from the growth rate in the Salomon Smith Barney Report (Mahedy et al. 2000).